arXiv:1704.01459v1 [physics.data-an] 5 Apr 2017

Statistical properties of interaction parameter estimates in direct coupling analysis

Yingying Xu
*Department of Computer Science, School of Science, Aalto University, P.O.Box 15400, FI-00076 Aalto, FINLAND**

Erik Aurell
Department of Computational Science and Technology, KTH-Royal Institute of Technology, SE-100 44 Stockholm, Sweden Depts of Applied Physics and Computer Science, Aalto University, FIN-00076 Aalto, Finland and Institute of Theoretical Physics, Chinese Academy of Sciences, Beijing, China

Jukka Corander
Department of Mathematics and Statistics, University of Helsinki, 00014 Helsinki, Finland; Department of Biostatistics, University of Oslo, 0317 Oslo, Norway

Yoshiyuki Kabashima
Department of Mathematical and Computing Science, School of Computing, Tokyo Institute of Technology, Tokyo 152-8552, Japan

We consider the statistical properties of interaction parameter estimates obtained by the direct coupling analysis (DCA) approach to learning interactions from large data sets. Assuming that the data are generated from a random background distribution, we determine the distribution of inferred interactions. Two inference methods are considered: the L_2 regularized naive mean-field inference procedure (regularized least squares, RLS), and the pseudo-likelihood maximization (plmDCA). For RLS we also study a model where the data matrix elements are real numbers, identically and independently generated from a Gaussian distribution; in this setting we analytically find that the distribution of the inferred interactions is Gaussian. For data of Boolean type, more realistic in practice, the inferred interactions do not generally follow a Gaussian. However, extensive numerical simulations indicate that their distribution can be characterized by a single function determined by a few system parameters after normalization by the standard deviation. This property holds for both RLS and plmDCA and may be exploitable for inferring the distribution of extremely large interactions from simulations for smaller system sizes.

Introduction—Identifying meaningful pairwise relationships between entities from very high-dimensional data is a central task in data science [1]. Of particular interest is the family of methods now generally known as direct coupling analysis (DCA), which aims to characterize such relationship by the parameters of a Potts model or an Ising model inferred from the data. Important progress has been made on inferring residue-residue contacts in proteins from multiple sequence alignments (MSA) of many homologous proteins [2–4], and has led to a breakthrough in *in silico* protein structure prediction [5, 6]. Among other applications we point to a method recently developed by us to detect of epistatic interactions in bacteria from population-wide dense sequencing data [7]. The area was reviewed from the methodological point of view at an early stage in [8] and more recently in [9] and [10].

The above mentioned methods and applications share several challenging aspects. The first, discussed in detail in [10] and [11] (and elsewhere), is that learning non-trivial probabilistic models from large data sets is computationally hard; in practice it can only be done approximately or relying on learning criteria weaker than maximum likelihood. The second is that the relevant data typically have fewer samples (n) than the number of parameters characterizing the model (p). For the protein contact prediction problems n is often in the range $10^3 - 10^5$ while the number of parameters in the corresponding Potts models grows quadratically with the length of the protein, and can hence realistically be $10^6 - 10^7$. Inference schemes therefore need to be regularized, as will be exemplified below by the RLS and plmDCA algorithms. The third aspect is that while at an intermediate stage one

* Email address: yingying.xu@aalto.fi

infers many more parameters than there is data, the final goal is to provide only (relatively) a small number k of predictions: in all successful applications to date $k \ll p$, and usually k has also been smaller (or much smaller) than n.

From all the above follows that the important task of assessing the statistical properties and significance thresholds of DCA is highly non-trivial. Indeed, in applications one would like to assign a p-value of a list of, say, k largest predictions. Alternatively, given an MSA representing some data one would like to be able to say that k pairwise relationships can be predicted from this data with some pre-determined level of confidence. Usually neither of the above has been done, but instead inference algorithms have been evaluated by comparing to a "ground truth" deduced from separate experimental data, in particular, for the contact prediction problem, by comparing to protein crystal structures [2, 4, 9, 12–14]. For a majority of potential future applications such a ground truth will not be available and there is therefore a need to develop statistical tests of the results of the inference. In fact the only previous step in this direction that we are aware of is our recent use of extreme value distribution theory (Gumbel distribution) to describe the background distribution for a variant of DCA where only the exceptionally large sampled predictions are retained in an intermediate step [7], and which is therefore a kind of special case.

In this paper we introduce the systematic study of the background and statistical significance of DCA predictions. We point out that the question of statistical significance of DCA is a large deviation problem for a nonlinear transformation of a data matrix, and discuss several examples which can be understood numerically and/or analytically.

Naive mean-field inference of Ising models and Regularized Least Squares— Consider a data matrix consisting of Boolean variables $\sigma_i^{(r)}$ taking values ± 1 and let the inference task be to estimate the interaction parameters ("J") in an Ising model

$$P(\underline{\boldsymbol{\sigma}}) = \frac{1}{Z} \exp\left(\sum_i h_i \sigma_i + \sum_{i \neq j} J_{ij} \sigma_i \sigma_j\right) \qquad (1)$$

from this data. The simplest version of variational inference is "naive mean-field" [1] where $\boldsymbol{J} = -\boldsymbol{C}^{-1}$ and $C_{ij} = \langle \sigma_i \sigma_j \rangle - \langle \sigma_i \rangle \langle \sigma_j \rangle$ is the covariance matrix. Note that the indices of J_{ij} are defined to be different and that therefore only the off-diagonal elements of the inverse covariance matrix $\boldsymbol{C}^{-1}$ are used. When $n < L$ naive mean-field is not a well-defined procedure since the covariance matrix then does not have full rank, and hence has to be regularized by the use of pseudo-counts, as in [4], or a sparsity-promoting L_1 regularizer, as in [15].

In the same family Regularized Least Squares (RLS) is an L_2-regularized inference method and given by the simple matrix equation

$$\boldsymbol{J}^{\mathrm{RLS}} = -\boldsymbol{C}\left(\eta \mathbf{1} + \boldsymbol{C}^2\right)^{-1} \qquad (2)$$

where $\mathbf{1}$ is the identity matrix and $\eta > 0$ is a positive regularization parameter. Although arguably one of the simplest non-trivial inference procedures one can imagine, to the best of our knowledge RLS was only introduced in the context of DCA in [14], and has not been investigated since in the subsequent DCA literature. We will use it here since the matrix form Eq. (2) renders it quite convenient for our purposes.

A solvable example— A model problem of RLS inference can be completely understood by random matrix theory. Assume that the elements in the data matrix $X \in \mathbb{R}^{n \times L}$ are real and Gaussian distributed $\mathcal{N}(0, \sigma^2)$. The covariance matrix $\boldsymbol{C} = \frac{1}{n} X^\top X$ is then a Wishart matrix. The spectrum of $\boldsymbol{C}$ converges by the Marcenko-Pastur law almost surely to a limit when n and L tend simultaneously to infinity, and since $\boldsymbol{J}^{\mathrm{RLS}}$ is related to $\boldsymbol{C}$ by (2) the spectrum of $\boldsymbol{J}^{\mathrm{RLS}}$ is almost surely a non-linear transformation of the Marcenko-Pastur distribution. Furthermore, the distribution of the individual elements of the distribution of the individual elements of $\boldsymbol{J}^{\mathrm{RLS}}$ can be obtained from a singular value decomposition of matrix X,

$$X = USV^\top, \qquad (3)$$

and regarding the left and right eigen-bases as samples from the uniform distributions of orthogonal matrices. The covariance matrix can then be expressed as

$$\begin{aligned} \boldsymbol{C} &= \frac{1}{n} V S^2 V^\top \\ &= V \Lambda V^\top \\ &= \left(\sum_{k=1}^{L} u_{ik} \lambda_k u_{jk}\right), \end{aligned} \qquad (4)$$

where λ_k is the kth eigenvalue of covariance matrix $\boldsymbol{C}$ and u_{ik} is the kth elements of ith eigen-base of matrix $\boldsymbol{C}$. From (2) the inferred interaction matrix obtained by RLS is

$$\left(J_{ij}^{\mathrm{RLS}}\right) = \left(\sum_{k=1}^{L} \frac{\lambda_k}{\eta + \lambda_k^2} u_{ik} u_{jk}\right). \qquad (5)$$

In this situation u_{ik} are samples from the uniform distributions of orthogonal matrices, and when the dimension of the matrix $\boldsymbol{C}$ goes to infinity u_{ik} can be

handled as random numbers [16] that satisfy

$$\overline{u_{ik}} = 0, \overline{u_{ik}u_{jl}} = \frac{1}{L}\delta_{ij}\delta_{kl}. \quad (6)$$

Condition (6) means that each diagonal component converges to an $\mathcal{O}(1)$ constant as

$$J_{ii}^{\mathrm{RLS}} = \langle \frac{\lambda}{\eta + \lambda^2} \rangle, \quad (7)$$

where the brackets $\langle \cdot \rangle$ denote the expectation with respect to eigenvalue distribution $\rho(\lambda)$ of the covariance matrix $\boldsymbol{C}$. The off-diagonal elements similarly follow a zero mean Gaussian distribution

$$J_{ij}^{\mathrm{RLS}} \sim \mathcal{N}(0, v_J^2) \quad (8)$$

with variance

$$v_J^2 \simeq \frac{1}{L-1}\left(\langle \left(\frac{\lambda}{\eta + \lambda^2} \right)^2 \rangle - \langle \frac{\lambda}{\eta + \lambda^2} \rangle^2 \right). \quad (9)$$

To test the above theory we generated standard i.i.d Gaussian random variables as data matrix elements, representative results summarized in Fig. 1. This indicates that after standardization $J_{ij} \leftarrow (J_{ij} - \overline{J})/v_J$, where the mean $\overline{J}$ vanishes in the current case, the coupling distribution collapses to a *single function* of $\mathcal{N}(0,1)$ irrespectively of any system parameters.

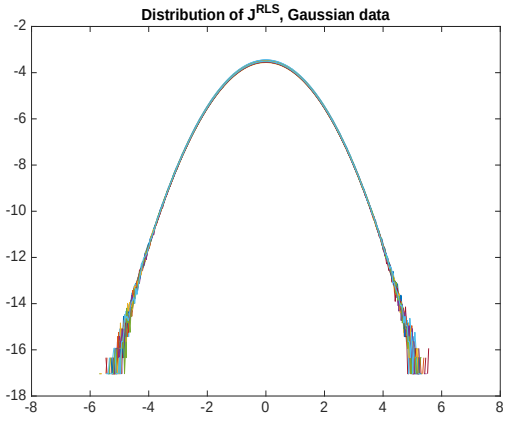

FIG. 1. Log-histogram of inferred off-diagonal interactions $\boldsymbol{J}^{\mathrm{RLS}}$ using RLS algorithm (Eq. (2) starting from $\mathcal{N}(0,1)$ i.i.d. Gaussian data, $\eta = 0.1$. Number of loci in this series was $L = 10000$ and number of samples $n = 1000, 2000, \ldots, 20000$. All distributions scaled by the standard deviation of each instance.

Agreement between numerics and the theoretical prediction is good. However, as discussed above DCA works by selecting only the exceptionally large predictions and what would matter for an evaluation of the statistical background is the distribution over such rare events. The deviation in the tail (or any finite range) between an empirical and a predicted probability distribution can be quantified by *e.g.* an Anderson-Darling test [17], but to connect to current practice in the DCA field we instead display in Fig. 2 the largest predicted couplings in a rank plot. Clearly the empirical distribution is quite close to Gaussian also according to this more challenging test.

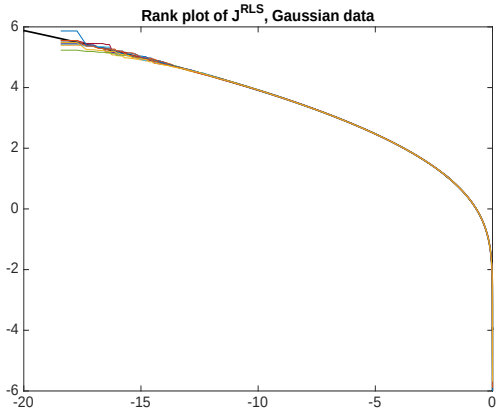

FIG. 2. Rank plot of inferred off-diagonal couplings. The J's are centered and scaled by the standard deviation of each instance, and then re-ordered as $J^1 \geq J^2 \geq \cdots J^r \geq \cdots$ where r indicates the rank. Data lines (colour online): J^r vs. $\log(r/N_L)$ where $N_L = L(L-1)$ is the total number of inferred off-diagonal couplings. Black curve: $x = \log \frac{1}{\sqrt{2\pi}} \int_y^\infty e^{-\frac{1}{2}t^2}\, dt$.

Boolean data, plmDCA: finite-size scaling property— We turn to RLS-inferred couplings from random Boolean data. As in realistic examples the bias of each variable in each position i varies, we examined the case where the bias $f_i = \Pr(\sigma_i = +1) = 1 - \Pr(\sigma_i = -1)$ is uniformly distributed on region $(0,1)$. The obtained coupling distribution is neither a Gaussian nor another analytically expressible distribution. Nevertheless, the experimental results indicate that as L grows keeping aspect ratio $\alpha = n/L$ and η fixed, the coupling distribution after standardization collapses to a single function characterized by α and η. In analogy with well-established use in percolation theory and constraint satisfaction we term this a *finite-size scaling property* [18].

This may be exploitable for practical purposes. As RLS needs the matrix inversion operation, experimental evaluation of the background distribution becomes

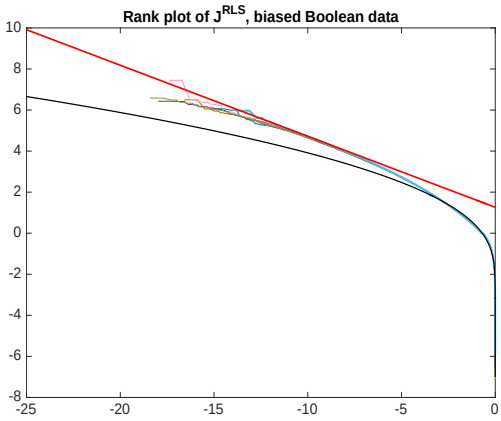

FIG. 3. Rank plot of RLS algorithm inferred off-diagonal couplings from random Boolean data generated as the bias of each variable in each position f_i were uniformly distributed on region $(0, 1)$. In this series, the aspect ratio $n/L = 0.1$ and the regularization parameter $\eta = 0.5$. The thick light blue curve was the average of 100 experiments for random data of loci size $L = 1000$, which is a smooth curve. The straight line was extrapolated by the very tail 100 points (about top 0.01% rank data) of this light blue curve. Number of loci in the colour curves which has a vibrating tail were $L = 2000, 4000, \ldots, 10000$. The black curve is drawn by standard Gaussian as a refernce line. All curves scaled by the standard deviation of each instance.

computationally infeasible as L very large. However, the scaling property allows us to infer the distribution of large L via Monte Carlo assessments for smaller L. A crucial drawback of this approach is that the information for the extremely large predictions that occur with a probability smaller than $2/(L(L-1))$ cannot be evaluated accurately. However, the empirical observation that the tails of the normalized coupling distribution typically go down like a Gaussian (Fig.3) indicates that the straight line extrapolated from the very tail in the rank plot acts as an upper bound for the extremely rare predictions, which can be useful for screening relevant couplings from the background.

We have thus far analyzed the RLS algorithm, which is based on mean field assumption, because of the simplicity of the resulting mathematical and numerical analysis. Another approximation algorithm for estimating the couplings in (1) is pseudo-likelihood maximization (plmDCA) [13], which has been shown to have very accurate performance in predicting the residue contacts in protein structures. plmDCA maximizes probability of each element σ_i conditioned by all other elements $\boldsymbol{\sigma}_{\setminus i}$, $P(\sigma_i|\boldsymbol{\sigma}_{\setminus i})$, and regularizes the couplings $\boldsymbol{J}$ and external fields $\boldsymbol{h}$ by the L_2 norm. For details, see Ref. [13]. Here we examine the coupling

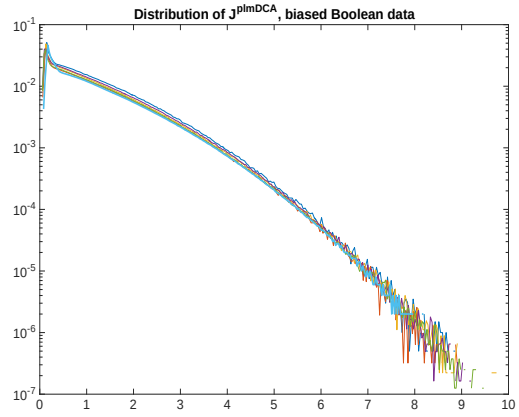

FIG. 4. Log-histogram of plmDCA algorithm inferred off-diagonal interactions from random Boolean data generated as f_i were uniformly distributed on region $(0, 1)$. In the experiments, $n/L = 0.1$ and $\lambda_J = \lambda_h = 0.5$. The thick light blue curve was the average of 100 experiments for random data of loci size $L = 1000$, which is a smooth curve in low ranking region. Number of loci in the colour curves which has a vibrating tail were $L = 2000, 2500, \ldots, 4000$. All distributions scaled by the standard deviation of each instance.

distribution for plmDCA when the input Boolean data follows the uniform bias distribution. The result is shown in Fig.4. Since Frobenius norm was taken on $||\boldsymbol{J}_{ij}(a, b)||_F$ for each i, j pairs, the output coupling value only has positive values.

Unlike RLS, the relation between the coupling obtained by plmDCA, $\boldsymbol{J}^{\mathrm{plmDCA}}$, and $\boldsymbol{C}$ is highly nontrivial, and cannot be analytically expressed. However, the experimental results (Fig.5) show that the scaling property holds for plmDCA as well, and can be used for inferring the distribution of extremely large couplings from extensive simulations for smaller system sizes. The value of this property may be more significant for plmDCA than for RLS since plmDCA is computationally demanding and practically difficult to repeat random simulations of large systems many times for accurately assessing the background effect.

Discussion—The study of the backgound distribution of inferred DCA couplings is of importance for supporting the selection of significant predictions. Our work in this paper proves that when the data matrix elements are i.i.d Gaussian variables the distribution of the inferred interactions is also Gaussian. For data of Boolean type, extensive numerical simulations indicate that their distribution can be characterized by a single function determined by a few

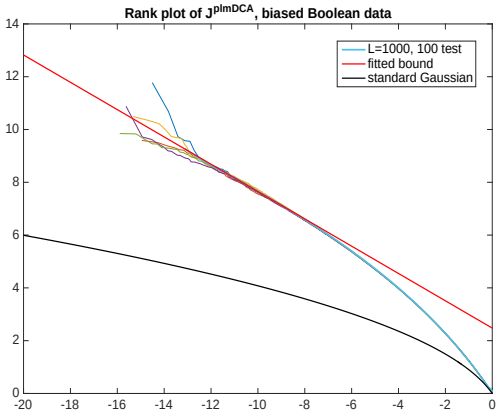

FIG. 5. Rank plot of plmDCA algorithm inferred off-diagonal couplings from random Boolean data same as in Fig.4 with the same parameter settings. The thick light blue curve was the average of 100 experiments for random data of loci size $L = 1000$, which is a smooth curve. The straight line was extrapolated by the very tail 100 points of this light blue curve. Other single test with different size L are fluctuated around the averaged curve in the tail. Black curve: standard Gaussian. All curves scaled by the standard deviation of each instance.

system parameters (aspect ratio α, regularization parameter and column bias $\{f_i\}$) after normalization by the standard deviation; this holds for both RLS and plmDCA. This property may be exploitable for inferring the distribution of extremely large interactions from simulations for smaller system sizes.

Also, investigations on other background models, for example, a neutral model, could be helpful to separate clonal inheritance[19] from epistasis[20] as a mechanism for linkage disequilibrium[21]. In our recent application of DCA to detect epistasis from genome sequence data[7], taking such biological effects into account in the background model could lead to a more accurate method to predict significant couplings.

I. ACKNOWLEDGEMENTS

This research is supported by the Swedish Science Council through grant 621-2012-2982(EA) and by the Academy of Finland through its Center of Excellence COIN (YX, EA and JC), and JSPS KAKENHI Nos. 25120013 and 17H00764 (YK). EA acknowledges support from Chinese Academy of Sciences CAS President's International Fellowship Initiative (PIFI) GRANT No. 2016VMA002.

[1] M. J. Wainwright and M. I. Jordan, Foundations and Trends in Machine Learning **1**, 1 (2008).
[2] M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, and T. Hwa, Proc. Natl. Acad. Sci. U. S. A.; arXiv:0901.1248v1 **106**, 67 (2009).
[3] L. Burger and E. v. Nimwegen, PLoS Computational Biology **6** (2010).
[4] F. Morcos, A. Pagnani, B. Lunt, A. Bertolino, D. S. Marks, C. Sander, R. Zecchina, J. N. Onuchic, T. Hwa, and M. Weigt, Proc. Natl. Acad. Sci. **108**, E1293 (2011).
[5] D. S. Marks, T. A. Hopf, and C. Sander, Nature Biotechnology **30**, 1072 (2012).
[6] S. Ovchinnikov, H. Park, N. Varghese, P.-S. Huang, G. A. Pavlopoulos, D. E. Kim, H. Kamisetty, N. C. Kyrpides, and D. Baker, **355**, 294 (2017).
[7] M. J. Skwark, N. J. Croucher, S. Puranen, C. Chewapreecha, M. Pesonen, Y. Y. Xu, P. Turner, S. R. Harris, S. B. Beres, J. M. Musser, J. Parkhill, S. D. Bentley, E. Aurell, and J. Corander, PLoS Genetics (2017), 10.1371/journal.pgen.1006508.
[8] Y. Roudi, E. Aurell, and J. Hertz, Front. Comput. Neurosci. **3** (2009).
[9] R. R. Stein, D. S. Marks, and C. Sander, PLoS Comput Biol **11**, e1004182 (2015).
[10] H. C. Nguyen, R. Zecchina, and J. Berg, "Inverse statistical problems: from the inverse ising problem to data science," arXiv:1702.01522 (2017).
[11] A. Montanari, Electron. J. Statist. **9**, 2370 (2015).
[12] M. Ekeberg, C. Lövkvist, Y. Lan, M. Weigt, and E. Aurell, Physical Review E **87** (2013).
[13] M. Ekeberg, T. Hartonen, and E. Aurell, Journal of Computational Physics **276**, 341 (2014).
[14] M. Andreatta, S. Laplagne, S. C. Li, and S. Smale, arXiv **1311.1301v3** (2014).
[15] D. T. Jones, D. W. A. Buchan, D. Cozzetto, and M. Pontil, Bioinformatics **28**, 184 (2012).
[16] K. Takeda, S. Uda, and Y. Kabashima, Europhysics letters **76**, 1193 (2006).
[17] T. W. Anderson and D. A. Darling, Annals of Mathematical Statistics **23**, 193 (1952).
[18] S. Kirkpatrick and B. Selman, **264**, 1297 (1994).
[19] R. Stein, Y. Gruenbaum, Y. Pollack, A. Razin, and H. Cedar, Proc Natl Acad Sci U S A **79**, 61 (1982).
[20] H. J. Cordell, Human Molecular Genetics **11**, 2463 (2002).
[21] M. Slatkin, Nat Rev Genet **9**, 477 (2008).